

WHAT IS CLAIMED IS:

- 1 1. For use with a multi-stage switch having  
2 - a first number,  $k \times n$ , of output ports,  
3 - a plurality of central modules, each having  
4 outgoing links, and  
5 - a second number of input modules, each including  $k$   
6 groups of  $n$  virtual output queues and outgoing links  
7 coupled with each of the plurality of central modules,  
8 and  
9 - a third number of sub-schedulers, each of the third  
10 number of sub-schedulers being able to arbitrate  
11 matching an input port with an outgoing link of one of  
12 the plurality of central modules via an outgoing link  
13 of the input module including the input port,  
14 a method for scheduling the dispatch of cells stored in the  
15 virtual output queues, the method comprising for each of  
16 the sub-schedulers, performing a matching operation, if it  
17 has been reserved, to match a cell buffered at a virtual  
18 output queue with an outgoing link of one of the plurality  
19 of central modules, wherein the matching operation  
20 includes:  
21 a) matching a non-empty virtual output queue of an  
22 input module with an outgoing link in the input  
23 module, wherein the outgoing link has an associated  
24 master arbitration operation for selecting one of the  
25  $k$  groups of  $n$  virtual output queues; and  
26 b) matching the outgoing link with an outgoing link  
27 of one of the plurality of central modules,  
28 wherein each of the sub-schedulers requires more  
29 than one cell time slot to generate a match from its  
30 matching operation, and

31            wherein the sub-schedulers can collectively  
32 generate a match result in each cell time slot.

1    2. The method of claim 1 wherein the act of matching a  
2 non-empty virtual output queue of an input module with an  
3 outgoing link in the input module includes:

4            i) sending, on behalf of each non-empty virtual  
5 output queue, a request to slave arbiters, each  
6 of the slave arbiters being associated with one  
7 of each of the outgoing links of the input  
8 module, and each of the slave arbiters being  
9 associated with the group of virtual output  
10 queues to which the non-empty virtual output  
11 queue belongs;

12           ii) sending, on behalf of each group of virtual  
13 output queues to which a non-empty virtual output  
14 queue belongs, a request to master arbiters, each  
15 of the master arbiters being associated with one  
16 of each of the outgoing links of the input  
17 module;

18           iii) selecting, with each of the master  
19 arbiters, a virtual output queue group having at  
20 least one non-empty virtual output queue, from  
21 among one or more virtual output queue groups  
22 that sent a request;

23           iv) selecting, with each of the slave arbiters,  
24 a non-empty virtual output queue, belonging to  
25 its associated group, from among one or more  
26 virtual output queues that sent a request; and

27           v) selecting, with the arbiter of the each of  
28 the selected non-empty virtual output queues of  
29 each of the selected virtual output queue groups,

30 an outgoing link from among the one or more  
31 candidate outgoing links, each of the one or more  
32 candidate outgoing links being associated with a  
33 master arbiter that selected the virtual output  
34 queue group and a slave arbiter that selected the  
35 non-empty virtual output queue.

1 3. The method of claim 2 wherein an act of selecting, with  
2 a master arbiter, a virtual output queue group having at  
3 least one non-empty virtual output queue, is done in  
4 accordance with a round robin discipline.

1 4. The method of claim 2 wherein an act of selecting, with  
2 a slave arbiter, a non-empty virtual output queue,  
3 belonging to its associated group, is done in accordance  
4 with a round robin discipline.

1 5. The method of claim 2 wherein the act of selecting,  
2 with the arbiter of the each of the selected non-empty  
3 virtual output queues of each of the selected virtual  
4 output queue groups, an outgoing link from among the one or  
5 more candidate outgoing links, is done in accordance with a  
6 round robin discipline.

1 6. The method of claim 2 wherein the acts of  
2 i) sending, on behalf of each non-empty virtual  
3 output queue, a request to slave arbiters, each  
4 of the slave arbiters being associated with one  
5 of each of the outgoing links of the input  
6 module, and each of the slave arbiters being  
7 associated with the group of virtual output

TELETYPE

8 queues to which the non-empty virtual output  
 9 queue belongs;  
 10 ii) sending, on behalf of each group of virtual  
 11 output queues to which a non-empty virtual output  
 12 queue belongs, a request to master arbiters, each  
 13 of the master arbiters being associated with one  
 14 of each of the outgoing links of the input  
 15 module;  
 16 iii) selecting, with each of the master  
 17 arbiters, a virtual output queue group having at  
 18 least one non-empty virtual output queue, from  
 19 among one or more virtual output queue groups  
 20 that sent a request;  
 21 iv) selecting, with each of the slave arbiters,  
 22 a non-empty virtual output queue, belonging to  
 23 its associated group, from among one or more  
 24 virtual output queues that sent a request; and  
 25 v) selecting, with the arbiter of the each of  
 26 the selected non-empty virtual output queues of  
 27 each of the selected virtual output queue groups,  
 28 an outgoing link from among the one or more  
 29 candidate outgoing links, each of the one or more  
 30 candidate outgoing links being associated with a  
 31 master arbiter that selected the virtual output  
 32 queue group and a slave arbiter that selected the  
 33 non-empty virtual output queue,  
 34 are performed at least twice within the third number of  
 35 cell time slots.

1 7. The method of claim 1 wherein the act of matching the  
 2 outgoing link of the input module with an outgoing link of  
 3 one of the central modules includes:

4           i) sending a request for the outgoing link of  
5           the input module to an arbiter for each of the  
6           outgoing links of the central modules that leads  
7           towards an output port associated with the  
8           virtual output queue matched with the outgoing  
9           link of the input module; and  
10          ii) selecting with the arbiter of each of the  
11          outgoing links of the central modules, an  
12          outgoing link of an input module from among those  
13          that sent a request.

1   8. The method of claim 7 wherein the act of selecting with  
2   the arbiter of each of the outgoing links of the central  
3   module, an outgoing link of the input module that broadcast  
4   a request, is done based on a round robin discipline.

1   9. The method of claim 1 further comprising:  
2       c) if a cell buffered at a virtual output queue has  
3       been successfully matched with its corresponding  
4       output port, informing the virtual output queue.

1   10. The method of claim 9 further comprising:  
2       d) for each of the virtual output queues, if the  
3       virtual output queue has been informed that it has  
4       been successfully matched with its corresponding  
5       output port, then dispatching its head of line cell.

1   11. The method of claim 1 wherein each of the virtual  
2   output queues is associated with a first count for  
3   indicating whether the virtual output queue is storing a  
4   cell awaiting dispatch, wherein a first count is

5 incremented upon learning that a new cell has arrived at  
6 its associated virtual output queue.

1 12. The method of claim 11 wherein the count is  
2 decremented when an available sub-scheduler is reserved for  
3 considering a head of line cell at a corresponding virtual  
4 output queue.

1 13. The method of claim 1 further comprising:  
2 c) for each of the sub-schedulers, maintaining a  
3 second indicator for each of the virtual output  
4 queues, for indicating whether the sub-scheduler is  
5 available or reserved,  
6 wherein the second indicator, for each of the  
7 sub-schedulers, is set to indicate that the associated  
8 sub-scheduler is reserved if the first indicator indicates  
9 that a corresponding virtual output queue is storing a cell  
10 awaiting dispatch arbitration.

1 14. The method of claim 1 further comprising:  
2 c) for each of the sub-schedulers, maintaining a  
3 second indicator for each of the virtual output  
4 queues, for indicating whether the sub-scheduler is  
5 available or reserved,  
6 wherein the second indicator, for each of the  
7 sub-schedulers, is set to indicate that the associated  
8 sub-scheduler is available if the associated sub-scheduler  
9 matches a cell buffered at a virtual output queue with its  
10 corresponding output port.

1 15. The method of claim 1 further comprising:

2 c) for each of the sub-schedulers, maintaining a  
3 second indicator for each of the virtual output  
4 queues, for indicating whether the sub-scheduler is  
5 available or reserved,

6 wherein the second indicator is set to indicate  
7 that a  $p^{\text{th}}$  sub-scheduler is reserved if the first indicator  
8 indicates that a corresponding virtual output queue is  
9 storing a cell awaiting dispatch arbitration,

10 wherein  $p$  is set to the current cell time slot  
11 modulo the third number.

1 16. For use with a multi-stage switch including

2 - a plurality of central modules, each including  
3 outgoing links towards output modules, the output  
4 modules collectively including a first number,  $k \times n$ ,  
5 of output ports;

6 - a second number of input modules, each including  
7 -  $k$  groups of  $n$  virtual output queues, and  
8 - outgoing links coupled with each of the  
9 plurality of central modules;

10 a dispatch scheduler comprising:

11 a) a third number of sub-schedulers; and

12 b) a first indicator, associated with each of the  
13 virtual output queues, for indicating whether the  
14 virtual output queue is storing a cell awaiting  
15 dispatch arbitration,

16 wherein each of the sub-schedulers is adapted to  
17 perform a matching operation, if it has been reserved, to  
18 match a cell buffered at a virtual output queue with its  
19 corresponding output port, and includes:

20 i) master arbiters, each of the master arbiters  
21 being associated with one of the outgoing links  
22 of the input module, for selecting a group of  
23 virtual output queues from among those associated  
24 with a received request,  
25 ii) groups of slave arbiters, each group of  
26 slave arbiters being associated with one of the k  
27 groups of n virtual output queues, for selecting  
28 a virtual output queue from among those  
29 submitting a request, and  
30 iii) virtual output queue arbiters, each virtual  
31 output queue arbiter being associated with one of  
32 the virtual output queues, for selecting an  
33 outgoing link of the input module from among  
34 those submitting a grant  
35 wherein each of the sub-schedulers requires more  
36 than one cell time slot to generate a match from its  
37 matching operation, and  
38 wherein the sub-schedulers can collectively  
39 generate a match result in each cell time slot.

1 17. The dispatch scheduler of claim 16 wherein each of the  
2 sub-schedulers require no more than the third number of  
3 cell time slots to generate a match result from its  
4 matching operation.

1 18. The dispatch scheduler of claim 16 wherein if a cell  
2 buffered at a virtual output queue has been successfully  
3 matched with its corresponding output port, the virtual  
4 output queue is so informed.



1 19. The dispatch scheduler of claim 16 wherein if a cell  
2 buffered at a virtual output queue has been successfully  
3 matched with its corresponding output port, its head of  
4 line cell is dispatched.

1 20. The dispatch scheduler of claim 16 wherein the first  
2 indicator, for each of the virtual output queues, for  
3 indicating whether the virtual output queue is storing a  
4 cell awaiting dispatch arbitration, is a count, and  
5 wherein the count is incremented upon learning  
6 that a new cell has arrived at the virtual output queue.

1 21. The dispatch scheduler of claim 20 wherein the count  
2 is decremented when an available sub-scheduler is reserved  
3 for considering a head of line cell at a corresponding  
4 virtual output queue.

1 22. The dispatch scheduler of claim 16 further comprising:  
2 c) a second indicator for each of the virtual output  
3 queues and for each of the sub-schedulers, indicating  
4 whether the sub-scheduler is available or reserved,  
5 wherein the second indicator, for each of the  
6 sub-schedulers, is set to indicate that the associated  
7 sub-scheduler is reserved if the first indicator indicates  
8 that a corresponding virtual output queue is storing a cell  
9 awaiting dispatch arbitration.

1 23. The dispatch scheduler of claim 16 further comprising:  
2 c) a second indicator for each of the virtual output  
3 queues and for each of the sub-schedulers, indicating  
4 whether the sub-scheduler is available or reserved,

5            wherein the second indicator, for each of the  
6 sub-schedulers, is set to indicate that the associated  
7 sub-scheduler is available if the associated sub-scheduler  
8 matches a cell buffered at a virtual output queue with its  
9 corresponding output port.

1    24. The dispatch scheduler of claim 16 further comprising:  
2        c) a second indicator for each of the virtual output  
3        queues and for each of the sub-schedulers, indicating  
4        whether the sub-scheduler is available or reserved,  
5            wherein the second indicator is set to indicate  
6        that a p<sup>th</sup> sub-scheduler is reserved if the first indicator  
7        indicates that a corresponding virtual output queue is  
8        storing a cell awaiting dispatch, and  
9            wherein p is set to the current cell time slot  
10       modulo the third number.

1    25. The dispatch scheduler of claim 16 wherein the means  
2       for matching a non-empty virtual output queue of an input  
3       module with an outgoing link in the input module further  
4       include:  
5            iv) means for sending, on behalf of each  
6            non-empty virtual output queue, a request to  
7            slave arbiters, each of the slave arbiters being  
8            associated with one of the outgoing links of the  
9            input module, and each of the slave arbiters  
10           being associated with one of the groups of  
11           virtual output queues; and  
12           v) means for sending, on behalf of each of the  
13           groups of virtual output queues to which a  
14           non-empty virtual output queue belongs, a request  
15           to master arbiters, each of the master arbiters

16                   being associated with one of the outgoing links  
17                   of the input module.

1   26.   The dispatch scheduler of claim 16 wherein each of the  
2   master arbiters operates in accordance with a round robin  
3   discipline.

1   27.   The dispatch scheduler of claim 26 wherein each of the  
2   master arbiters operates independent of the others.

1   28.   The dispatch scheduler of claim 16 wherein each of the  
2   slave arbiters operates in accordance with a round robin  
3   discipline.

1   29.   The dispatch scheduler of claim 28 wherein each of the  
2   slave arbiters operates independent of the others.

1   30.   The dispatch scheduler of claim 16 wherein each of the  
2   virtual output queue arbiters operates in accordance with a  
3   round robin discipline.

1   31.   The dispatch scheduler of claim 30 wherein each of the  
2   virtual output queue arbiters operates independent of the  
3   others.

1   32.   The dispatch scheduler of claim 16 wherein the means  
2   for matching a non-empty virtual output queue of the input  
3   module with an outgoing link in the input module performs  
4   multiple matching iterations within the third number of  
5   cell time slots.

